ε-Net Approach to Sensor k-Coverage

Giordano Fusco and Himanshu Gupta

Computer Science Department
Stony Brook University
Stony Brook, NY 11790
{fusco,hgupta}@cs.sunysb.edu

Abstract. Wireless sensors rely on battery power, and in many applications it is difficult or prohibitive to replace them. Hence, in order to prolongate the system's lifetime, some sensors can be kept inactive while others perform all the tasks. In this paper, we study the k-coverage problem of activating the minimum number of sensors to ensure that every point in the area is covered by at least k sensors. This ensures higher fault tolerance, robustness, and improves many operations, among which position detection and intrusion detection.

The k-coverage problem is trivially NP-complete, and hence we can only provide approximation algorithms. In this paper, we present an algorithm based on an extension of the classical ε-net technique. This method gives a $O(\log M)$-approximation, where M is the number of sensors in an optimal solution. We do not make any particular assumption on the shape of the areas covered by each sensor, besides that they must be closed, connected and without holes.

1 Introduction

Coverage problems have been extensively studied in the context of sensor networks (see for example [1–4]). The objective of sensor coverage problems is to minimize the number of active sensors, to conserve energy usage, while ensuring that the required region is sufficiently monitored by the active sensors. In an over-deployed network we can also seek k-coverage, in which every point in the area is covered by at least k sensors. This ensures higher fault tolerance, robustness, and improves many operations, among which position detection and intrusion detection.

The k-coverage problem is trivially NP-complete, and hence we focus on designing approximation algorithms. In this paper, we extend the well-known ε-net technique to our problem, and present an $O(\log M)$-factor approximation algorithm, where M is the size of the optimal solution. The classical greedy algorithm for set cover [5], when applied to k-coverage, delivers a $O(k \log n)$-approximation solution, where n is the number of *target points* to be covered. Our approximation algorithm is an improvement over the greedy algorithm, since our approximation factor of $O(\log M)$ is independent of k and of the number of target points.

Instead of solving the sensor's k-coverage problem directly, we consider a dual problem, the k-hitting set. In the k-hitting set problem, we are given sets and points, and we look for the minimum number of points that "hit" each set at least k times (a set is hit by a point if it contains it). Brönnimann and Goodrich were the first [6] to solve the hitting set using the ε-net technique [7]. In this paper, we introduce a generalization of ε-nets, which we call (k, ε)-nets. Using (k, ε)-nets with the Brönnimann and Goodrich algorithm's [6], we can solve the k-hitting set, and hence the sensor's k-coverage problem. Our main contribution is a way of constructing (k, ε)-nets by random sampling. A recent Infocom paper [8] uses ε-nets to solve the k-coverage problem. However we believe that their result is fundamentally flawed (see Section 2.1 for more details). So, to the best of our knowledge, we are the first to give a correct extension of ε-nets for the k-coverage problem.

(a) (b)

Fig. 1. Sensing regions associated with a sensor: (a) a general shape that is closed, connected, and without holes, and (b) a disk.

Paper Organization. The rest of the paper is organized as follow. The k-coverage problem is introduced in Section 2. Section 2.1 contains detailed discussion about related work. The ε-net approach is presented in Section 3.

2 Problem Formulation and Related Work

We start by defining the sensing region and then we will define the k-coverage problem with sensors. In the literature, sensing regions have been often modeled as disks. In this paper, we consider sensing regions of general shape, because this reflects a more realistic scenario.

Definition 1 (Sensing Region). *The sensing region of a sensor is the area "covered" by a sensor. Sensing regions can have any shape that is closed, connected, and without holes, as in Fig. 1(a). Often, sensing regions are modeled as disks as in Fig. 1(b), but we consider more general shapes.*

Definition 2 (Target Points). Target points *are the given points in the 2D plane that we wish to cover using the sensors.*

k-SC Problem. Given a set of sensors with fixed positions and a set of target points, select the minimum number of sensors, such that each target point is covered (is contained in the selected sensing region) by at least k of the selected sensors.

For simplicity, we have defined the above k-SC problem's objective as coverage of a set of given target *points*. However, as discussed later, our algorithms and techniques easily generalize to the problem of covering a given area.

Example 1. Suppose we are given 4 sensors and 20 points as in Fig. 2(a), and we want to select the minimum number of sensors to 2-cover all points. In this particular example, 2 sensors are not enough to 2-cover all points. Instead, 3 sensors suffices, as shown in Fig. 2(b).

2.1 Related Work

In recent years, there has been a lot of research done [2, 3, 9, 1] to address the coverage problem in sensor networks. In particular, Slijepcevic and Potkonjak [3] design a centralized heuristic to select mutually exclusive sensor covers that independently cover the network region. In [2], Charkrabarty et al. investigate linear programming techniques to optimally place a set of sensors on a sensor field for a complete coverage of the field. In [10], Shakkottai et al. consider an unreliable sensor network, and derive necessary and sufficient conditions for the coverage of the region and connectivity of the network with high probability. In one of our

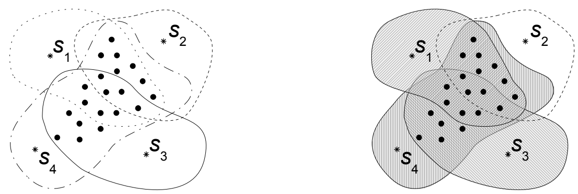

(a) (b)

Fig. 2. Illustrating k-SC problem. Suppose we are given 4 sensors of centers $s_1, \ldots, s_4$ and 20 points as in (a). The problem is to select the minimum number of sensors to k-cover all points. A possible solution for $k = 2$ is shown in (b), where 3 sensors suffice to 2-cover all points.

prior works [1], we designed a greedy approximation algorithm that delivers a connected sensor-cover within a logarithmic factor of the optimal solution; this work was later generalized to k-coverage in [11].

Recently, Hefeeda and Bagheri [8] used the well-known ε-net technique to solve the problem of k-covering the sensor's *locations*. However, we strongly believe that their result is fundamentally flawed.[1] In this article we present a correct extension of the ε-net technique for the k-coverage problem in sensor networks.

Two closely related problems to the sensor-coverage problem are set cover and hitting set problems. The area covered by a sensor can be thought as a set, which contains the points covered by that sensor. The hitting set problem is a "dual" of the set cover problem. In both set cover and hitting set problems, we are given sets and elements. While in set cover the goal is to select the minimum number of sets to cover all elements/points, in hitting set the goal is to select a subset of elements/points such that each set is hit. The classical result for set cover [5] gives a $O(\log n)$-approximation algorithm, where n is the number of *target points* to be covered. The same greedy algorithm also delivers a $O(k \log n)$-approximation solution for the k-SC problem. In contrast, the result in this article yields an $O(\log M)$-approximation algorithm for the k-SC problem, where M is the optimal size (i.e., minimum number of *sensors* needed to provide k-coverage of the given target points). Note that our approximation factor is independent of k and of the number of target points.

Brönnimann and Goodrich [6] were the first to use the ε-net technique [7] to solve the hitting set problem and hence the set cover with an $O(\log M)$-approximation, where M is the size of the optimal solution. In this article, we extend their ε-net technique to k-coverage. It is interesting to observe that our extension is independent of k and it gives a $O(\log M)$-approximation also for k-coverage. For the particular case of 1-coverage with disks, it is possible to build "small" ε-nets using the method of Matoušek, Seidel and Welzl [12], and obtain a constant-factor approximation for the 1-hitting set problem. Their method [12] can be easily extended to k-hitting set, and this would give a constant-factor approximation for the k-SC problem when the sensing regions are disks. However, in this paper we focus on sensing regions of arbitrary shapes and sizes, as long as they are closed, connected, and without holes.

Another related problem is the art gallery problem (see [13] for a survey) which is to *place* a minimum number of guards in a polygon so that each point in the polygon is visible from at least one of the guards. Guards may be looked upon as sensors with infinite range. However, in this paper, we focus on *selecting* already deployed sensor.

[1] Essentially, they select a set of *subsets* of size k (called k-flowers) represented by the center of their locations. However, their result is based on the following incorrect claim that if the centers of a set of k-flowers 1-covers a set of points N, then the set of sensors associated with the k-flowers will k-cover N. In addition, in their analysis, they implicitly assume that an optimal solution can be represented as a disjoint union of k-flowers, which is incorrect.

3 The ε-Net Based Approach

In this section, we present an algorithm based on the classical ε-net technique, to solve the k-coverage problem. The classical ε-net technique is used to solve the hitting set problem, which is the dual of the set cover problem. The k-SC problem is essentially a generalization of the set cover problem – thus, we will extend the ε-net technique to solve the corresponding generalization of the hitting set problem.

3.1 Hitting Set Problem and the ε-Net Technique

We start by describing the use of the classical ε-net technique to solve the traditional hitting set problem. We begin with a couple of formal definitions.

Set Cover (SC); Hitting Set (HS). Given a set of points X and a collection of sets $\mathcal{C}$, the *set cover* (SC) problem is to select the minimum number of sets from $\mathcal{C}$ whose union contains (covers) all points in X. The *hitting set* (HS) problem is to select the minimum number of points from X such that all sets in $\mathcal{C}$ are "hit" (a set is considered hit, if one of its points has been selected).

Note that HS is a *dual* of SC, and hence solving HS is sufficient to solve SC.

We now define ε-nets. Intuitively, an ε-net is a set of points that hits all *large* sets (but may not hit the smaller ones). For the overall scheme, we will assign weights to points, and use a generalized concept of weighted ε-nets that must hit all large-weighted sets.

Definition 3 (ε-Net; Weighted ε-Net). *Given a set system $(X, \mathcal{C})$, where X is a set of points and $\mathcal{C}$ is a collection of sets, a subset $H \in X$ is an* ε-net *if for every set S in $\mathcal{C}$ s.t. $|S| \geq \varepsilon|X|$, we have that $H \cap S \neq \emptyset$.*

Given a set system $(X, \mathcal{C})$, and a weight function $w : X \to Z^+$, define $w(S) = \sum_{x \in S} w(x)$ for $S \subseteq X$. A subset $H \subseteq X$ is a weighted ε-net *for $(X, \mathcal{C}, w)$ if for every set S in $\mathcal{C}$ s.t. $w(S) \geq \varepsilon \cdot w(X)$, we have that $H \cap S \neq \emptyset$.*

Using ε-Nets to Solve the Hitting Set Problem. The original algorithm for solving hitting set problem using ε-net was invented by Brönnimann and Goodrich [6]. Below, we give a high-level description of their overall approach (referred to as the BG algorithm), because it will help understand our own extension. We begin by showing how ε-nets are related to hitting sets, and then, show how to use ε-nets to actually compute hitting sets.

Let's assume that we have a black-box to compute weighted ε-nets, and that we know the optimal hitting set H^* which is of size M. Now, define a weight function w^* as $w^*(x) = 1$ if $x \in H^*$ and $w^*(x) = 0$ otherwise. Then, set $\varepsilon = 1/M$, and use the black-box to compute a weighted ε-net for $(X, \mathcal{C}, w^*)$. It is easy to see that this weighted ε-net is actually a hitting set for $(X, \mathcal{C})$, since $w^*(S) \geq \varepsilon w^*(X)$ for *all* sets $S \in \mathcal{C}$. There are known techniques [7] to compute weighted ε-nets of size $O((1/\varepsilon)\log(1/\varepsilon))$ for set systems with a constant VC-dimension (defined later); thus, the above gives us a $O(\log M)$-approximate solution. For the particular case of disks, it is possible to construct ε-nets of size $O(1/\varepsilon)$ [12], and hence obtain a constant-factor approximation.

However, in reality, we do not know the optimal hitting set. So, we iteratively guess its size M, starting with $M = 1$ and progressively doubling M until we obtain a hitting set solution (using the above approach). Also, to "converge" close to the w^* above, we use the following scheme. We start with all weights set to 1. If the computed weighted ε-net is not a hitting set, then we pick one set in $\mathcal{C}$ that is not hit by it and double the weights of all points that it contains. Then, we iterate with the new weights. It can be shown that if the estimate of M is correct and using $\epsilon = 1/(2M)$, then we are *guaranteed* to find a hitting set using the above approach after a certain number of iterations. Thus, if we don't find a hitting set after enough iterations, we double the estimate of M and try again. It can be shown [6] that the above approach finds an $O(\log M)$-approximate hitting set in polynomial time for set systems with constant VC-dimension (defined below), where M is the size of the optimal hitting set.

VC Dimension. We end the description of the BG algorithm, with the definition of Vapnik-Červonenkis (VC) dimension of set systems. Informally, the VC-dimension of a set system $(X, \mathcal{C})$ is a mathematical way of characterizing the "regularity" of the sets in $\mathcal{C}$ (with respect to the points X) in the system. A bounded VC-dimension allows the construction of an ε-net through random sampling of large enough size. The VC-dimension is formally defined in terms of set shattering, as follows.

Definition 4 (VC-Dimension). *A set S is considered to be* shattered *by a collection of sets $\mathcal{C}$ if for each $S' \subseteq S$, there exists a set $C \in \mathcal{C}$ such that $S \cap C = S'$. The* VC-dimension *of a set system $(X, \mathcal{C})$ is the cardinality of the largest set of points in X that can be shattered by $\mathcal{C}$.*

In our case, the VC dimension is at most 23 as given by the following theorem by Valtr [14].

Theorem 1. *If $X \subset \mathbb{R}^2$ is compact and simply connected, then VC-dimension of the set system $(X, \mathcal{C})$, where X is a set of points and $\mathcal{C}$ is a collection of sets, is at most 23.*

Note that for a finite collection of sensors, whose covering regions are compact and simply connected, the dual is compact and simply connected too.

3.2 k-Hitting Set Problem and The (k, ε)-net Technique

We now formulate the k-hitting set (k-HS) problem, which is a generalization of the hitting set problem, viz., we want each set in the system to be hit by k selected points.

Definition 5 (k-Hitting Set (k-HS)). *Given a set system $(X, \mathcal{C})$, the* k-hitting set (k-HS) *problem is to find the smallest subset of points $H \subseteq X$ with* at most one *point for each sibling-set such that H hits every set in $\mathcal{C}$ at least k times.*

Connection Between k-HS and k-SC Problem. Note that the above k-HS problem is the (generalized) dual of our sensor k-coverage problem (k-SC problem). Essentially, each point in the k-HS problem corresponds to a sensing region of a sensor, and each set in the k-HS problem corresponds to a target point. Below, we describe how to solve the k-HS problem, which essentially solves our k-SC problem. To solve the k-HS, we need to define and use a generalized notion of ε-net.

Definition 6 (Weighted (k, ε)-Net). *Suppose $(X, \mathcal{C})$ is a sibling-set system, and $w : X \rightarrow Z^+$ is a weight function. Define $w(S) = \sum_{x \in S} w(x)$ for $S \subseteq X$. A set $N \subseteq X$ is a* weighted (k, ε)-net *for $(X, \mathcal{C}, w)$ if $|N \cap S| \geq k$, whenever $S \in \mathcal{C}$ and $w(S) \geq \varepsilon \cdot w(X)$.*

Using (k, ε)-Nets to Solve k-HS. We can solve the k-HS problem using the BG algorithm [6], without much modification. However, we need an algorithm compute weighted (k, ε)-nets. The below theorem states that an appropriate random sampling of about $O(k/\varepsilon \log k/\varepsilon)$ points from X gives a (k, ε)-net with high probability, if the set system $(X, \mathcal{C})$ has a bounded VC-dimension. For the sake of clarity, we defer the proof of the following theorem.

Theorem 2. *Let $(X, \mathcal{C}, w)$ be a weighted set system. For a given number m, let $N(m)$ be a subset of points of size m picked randomly from X with probability proportional to the total weight of the points in such subset.*

Then, for

$$m \geq \max\left(\frac{2}{\varepsilon} \log_2 \frac{2}{\delta}, \ \frac{K}{\varepsilon} \log_2 \frac{K}{\varepsilon}\right) \tag{1}$$

the subset $N(m)$ is a weighted (k, ε)-sibling-net with probability at least $1 - \delta$, where $K = 4(d + 2k - 2)$, and d is the VC-dimension of the set system. □

Algorithm 1: Solving k-HS Problem using (k, ε)-nets.
Since k-HS is the dual of k-SC, this algorithm also solves k-SC
(in k-SC, X corresponds to the set of sensors, and $\mathcal{C}$ to the set of target points).

Given a set system $(X, \mathcal{C})$.
for $(M = 1;\ \ M <= |X|;\ \ M *= 2)$
  $\varepsilon = k/(2\,M)$;
  reset the weights of all points in X to 1;
  for $(i = 0;\ \ i < (4/k)M \log_2(|X|/M);\ \ i{+}{+})$
    Compute a (k, ε)-net N of size m using Theorem 2;
    if each set in $\mathcal{C}$ is k-hit by N, **return** N;
    select a set in $\mathcal{C}$ that is not k-hit, and double the weight of all the points in the set;

Now, based on the above theorem, we can use the BG algorithm with some modifications to solve the k-HS problem. Essentially, we estimate the size M of an optimal k-HS (starting with 1 and iteratively doubling it), set $\varepsilon = k/(2M)$, and use Theorem 2 to compute[2] a (k, ε)-net N of size m. If N is indeed a k-hitting set, we stop; else, we pick a set in the system $\mathcal{C}$ that is not k-hit and double the weight of all the points it contains. With the new weights, we iterate the process. It can be shown[3] that within $(4/k)M \log_2(n/M)$ iterations of weight-doubling, we are guaranteed to get a k-HS solution if the optimal size of a k-HS is indeed M. Thus, after $(4/k)M \log_2(n/M)$ iterations, if we haven't found a k-HS, we can double our current estimate of M, and iterate. See Algorithm 1. The below theorem shows that the above algorithm gives an $O(\log M)$-approximate solution in polynomial time with high probability for general sets. The proof of the following theorem is again similar to that for the BG algorithm [6].

Theorem 3. *The algorithm described above (Algorithm 1) runs in time* $O((|\mathcal{C}| + |X|)\,|X| \log |X|)$ *and gives a* $O(\log M)$*-approximate solution for the k-HS problem for a general set systems* $(X, \mathcal{C})$ *of constant VC-dimension, where* M *is the optimal size of a k-HS.*

Proof. The outer `for` loop, where M is doubled each time, is run at most $O(\log |X|)$ times. The inner `for` loop, where the weights are doubled for a set, is executed at most $\frac{4}{k} M \log_2 \frac{|X|}{M} = O(|X|/k)$ times. Computing a (k, ε)-net using Theorem 2 takes at most $O(|X|)$ time, while the doubling-weight process may take up to $O(|\mathcal{C}|)$ time.

We now prove the approximation factor. An optimal algorithm would find a k-hitting set of size M. If the VC-dimension is a constant, the k-HS method of Theorem 2 finds a (k, ε)-net of size $O(\frac{k}{\varepsilon} \log \frac{k}{\varepsilon})$. So if $\varepsilon = \frac{k}{2\,M}$, the size of the k-hitting set is $O(M \log M)$, which is a $O(\log M)$-approximation. □

Outline of Proof of Theorem 2. There are two challenges in generalizing the random-sampling technique of [7], viz., (i) sampling with replacement cannot be used, and (ii) weights must be part of the sampling process.

Challenges in Extending the Technique of [7] to k-hitting set. The classical method [7] of constructing an ε-net consists of randomly picking a set N of at least m points, for a certain m, where each point is picked *independently* and randomly from the given set of points. This way of constructing an ε-net may result in duplicate points in N, but *the presence of duplicates does not cause a problem in the analysis*. Thus, we can also construct weighted ε-net easily by emulating weights using duplicated copies of the same point. The above described approach works well for 1-hitting set, partly because we do not count the number of times each set is hit. However, for the case of k-hitting set, when constructing a (k, ε)-net, we need to ensure that

[2] Theorem 2 gives a (k, ε)-net with high probability. It is possible to check efficiently if the obtained set is indeed a (k, ε)-net. If it is not, we can try again until we get one. On average, a small number of trials are sufficient to obtain a (k, ε)-net.

[3] See Appendix A for the proof, which is similar to the one for BG in [6].

the number of *distinct* points that hit each set is at least k. Thus, constructing a (k, ε)-net by picking points independently at random (with duplicates) does not lead to correct analysis. Instead, we suggest a novel method to construct a weighted (k, ε)-net N by: (i) selecting a random subset of points (without duplicates) at once, and (ii) including the weights directly in the above sampling process. To the best of our knowledge, we are the first one to propose this extension.[4]

Proof Sketch of Theorem 2. Let m be as given by equation (1), and N be the subset of points randomly picked from X as described in Theorem 2. After picking N, pick another set T (for the purposes of the below analysis) in the same way as N. We now define two events

$$\begin{aligned} E_1 &= \{\exists A \in \mathcal{C} \text{ s.t. } |A \cap N| < k,\ w(A) \geq \varepsilon\, w(X)\} \\ E_2 &= \{\exists A \in \mathcal{C} \text{ s.t. } |A \cap N| < k,\ w(A) \geq \varepsilon\, w(X),\ |A \cap Z| \geq \varepsilon m\} \end{aligned}$$

where $Z = N \cup T$ and $\varepsilon m \leq \mathrm{E}[\,|A \cap Z|\,]/2$. The proof consists of 3 major steps:

1. First, we show that $\Pr[E_1] \leq 2 \Pr[E_2]$.
2. Then, it's easier to bound the probability of E_2
$$\Pr[E_2] \ \leq\ (2m)^{d+2k-2}\, 2^{-\varepsilon m}$$
3. Finally, we have that m verifies $2(2m)^{d+2k-2}\, 2^{-\varepsilon m} \ \leq\ \delta$

The outline of each step follows:

1. From the definition of conditional probability
$$\Pr[E_2 \mid E_1] \ =\ \Pr[E_2 \cap E_1]/\Pr[E_1] \ =\ \Pr[E_2]/\Pr[E_1]$$
So we just need to show that $\Pr[E_2 \mid E_1] \geq 1/2$. Let $Z = \{y_1, \ldots, y_{2m}\}$ (where y_i's *are* pairwise different). Define the random variable
$$Y_i \ =\ \begin{cases} 1 & \text{if } y_i \in A \\ 0 & \text{o.w.} \end{cases}$$
Set $Y = \sum_{i=1}^{2m} Y_i$, and we have $Y = |A \cap Z|$. It is possible to show that $\mathrm{E}[Y] = \mu \geq 2\varepsilon m$, and $\mathrm{Var}[Y] \leq \mu$. Applying Chebyshev's inequality
$$\Pr[Y < \frac{\mu}{2}] \ \leq\ \Pr[|Y - \mathrm{E}[Y]| > \frac{\mu}{2}] \ \leq\ \frac{4}{\mu^2}\,\mathrm{Var}[Y] \ \leq\ \frac{1}{2}$$
and the result follows.
2. We use an alternate view. Instead of picking N and then T, pick $Z \subseteq X$ of size $2m$, then pick $N \subseteq Z$ and set $T = Z \setminus N$. It can be shown that the two views are equivalent. Now, define
$$E_A \ =\ \{|A \cap N| < k,\ |A \cap Z| \geq \varepsilon m\}$$
Since N and T are disjoint $|A \cap Z| = |A \cap N| + |A \cap T|$ and then $|A \cap N| < k$ iff $|A \cap Z| < k + |A \cap T|$. So we have that E_A happens only if $\varepsilon m \leq |A \cap Z| \leq m + k - 1$. By counting the number of ways of choosing N s.t. $|A \cap N| < k$, we can bound $\Pr[E_A]$
$$\begin{aligned} \Pr[E_A] &= \Pr[|A \cap N| < k,\ |A \cap Z| \geq \varepsilon m] \\ &\leq \Pr[|A \cap N| < k \mid \varepsilon m \leq |A \cap Z| \leq m+k-1] \\ &\leq (2m)^{2k-2} \binom{2m-\varepsilon m}{m} \Big/ \binom{2m}{m} \ \leq\ (2m)^{2k-2}\, 2^{-\varepsilon m} \end{aligned}$$

[4] As discussed before, [8] uses ε-nets to solve sensor's k-coverage, but their method is flawed (see Footnote 1 for more details).

Since E_A depends only on the intersection $A \cap Z$

$$\Pr[E_2] \;\le \bigcup_{A \,|\, A\cap Z \text{ is unique}} \Pr[E_A] \;\le\; |\mathcal{C}_{|Z|}|\,(2m)^{2k-2}\,2^{-\varepsilon m}$$

3. Similar to [7]. □

Please, refer to Appendix B for the detailed proof.

Remark. Note that the approximation factor of Theorem 3 could be improved, if we could design an algorithm to construct smaller (k,ε)-nets. For instance, if we could construct a (k,ε)-net of size $O(k/\varepsilon)$, then we would have a constant-factor approximation for the k-HS problem. For the particular case of disks, it is easy to extend[5] the method in [12] to build a (k,ε)-net of size $O(k/\varepsilon)$ (see Appendix C for more details).

3.3 Distributed ε-Net Approach

Distributed implementation of the ε-net algorithm requires addressing the following main challenges.

1. We need to construct a (k,ε)-net, through some sort of distributed randomized selection.
2. For each constructed (k,ε)-net N, we need to verify in a distributed manner whether N is indeed a k-coverage set (k-hitting set in the dual).
3. If N is not a k-coverage set, then we need to select *one* target point (a set in the dual) that is not k-covered by N and double the weights of all the sensing regions covering it.

We address the above challenges in the following manner. First, we execute the distributed algorithm in *rounds*, where a round corresponds to one execution of the inner `for` loop of Algorithm 1 (i.e., execution of the sampling algorithm for a particular set of weights and a particular estimate of M). We implement rounds in a weakly synchronized manner using internal clocks. Now, for each of the above challenges, we use the following solutions.

1. Each sensor keeps an estimate of the total weight of the system, and computes m independently. To select m sensors, each sensor decides to select itself independently with a probability $p = m * own\_weight/total\_weight$, resulting in selection of m sensors (in expectation).
2. Locally, verify k-coverage of the owned target points, by exchanging messages with near-by (that cover a common target point) sensors. If a target point owned by a sensor D and its near-by sensors are all k-covered for a certain number of rounds (for example 10), then D exits the algorithm.
3. Each sensor decides to select one of the owned target points with a probability of $q = 1/((1-\varepsilon)\,n)$, which ensures that the expected number of selected target point is 1.

3.4 Generalizations to k-coverage of an Area

The ε-net approach can also be used to k-cover a given area, rather than a given set of target points (as required by the formulation of k-SC problem). Essentially, coverage of an area requires dividing the given area into "subregions" as in our previous work [1]; a subregion is defined as a set of points in the plane that are covered by the *same* set of sensing regions. The number of such subregions can be shown to be polynomial in the total number of sensing regions in the system. The algorithm described here can then be used without any other modification, and the performance guarantee still holds.

[5] Essentially, it is enough to replace $\delta = \varepsilon/6$ with $\delta = \varepsilon/(6k)$ and the proof follows through. Also note that the dual of disks and points is also composed by disks and points.

4 Conclusions

In this paper, we studied the k-coverage problem with sensors, which is to select the minimum number of sensors so that each target point is covered by at least k of them. We provided a $O(\log M)$-approximation, where M is the number of sensors in an optimal solution. We introduced a generalization of the classical ε-net technique, which we called (k, ε)-net. We gave a method to build (k, ε)-nets based on random sampling. We showed how to solve the sensor's k-coverage problem with the Brönnimann and Goodrich algorithm [6] together with our (k, ε)-nets. We believe to be the first one to propose this extension.

As a future work, we would like to extend this technique to *directional sensors*. A directional sensor is a sensor that has associated multiple sensing regions, and its *orientation* determines its actual sensing region. The k-coverage problem with directional sensors is NP-complete and in [15] we proposed a greedy approximation algorithm. We believe that the use of (k, ε)-nets can give a better approximation factor for this problem.

5 Acknowledgments

This work was supported in part by the NSF awards: 0713186, 0721701, 0721665.

References

1. Gupta, H., Zhou, Z., Das, S., Gu, Q.: Connected sensor cover: Self-organization of sensor networks for efficient query execution. ACM/IEEE Transactions on Networking (TON) **14**(1) (2006)
2. Charkrabarty, K., Iyengar, S., Qi, H., Cho, E.: Grid coverage for surveillance and target location in distributed sensor networks. IEEE Transaction on Computers (2002)
3. Slijepcevic, S., Potkonjak, M.: Power efficient organization of wireless sensor ad-hoc networks. In: Proceedings of the International Conveference of Communication (ICC). (2001)
4. Meguerdichian, S., Koushanfar, F., Qu, G., Potkonjak, M.: Exposure in wireless ad-hoc sensor networks. In: Proceedings of the International Conference on Mobile Computing and Networks (MobiCom). (2001)
5. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: Introduction to algorithms. 2nd edn. MIT Press McGraw-Hill, Cambridge Boston (2001)
6. Brönnimann, H., Goodrich, M.T.: Almost optimal set covers in finite vc-dimension. Discrete & Computational Geometry **14**(4) (1995) 463–479
7. Haussler, D., Welzl, E.: Epsilon-nets and simplex range queries. Discrete & Computational Geometry **2** (1987) 127–151
8. Hefeeda, M., Bagheri, M.: Randomized k-coverage algorithms for dense sensor networks. In: 26th IEEE International Conference on Computer Communications (INFOCOM). (2007) 2376–2380
9. Ye, F., Zhong, G., Lu, S., Zhang, L.: Peas: A robust energy conserving protocol for long-lived sensor networks. In: Proceedings of the International Conference on Distributed Computing Systems. (2003)
10. Shakkottai, S., Srikant, R., Shroff, N.: Unreliable sensor grids: Coverage, connectivity and diameter. In: Proceedings of the IEEE INFOCOM. (2003)
11. Zhou, Z., Das, S.R., Gupta, H.: Connected k-coverage problem in sensor networks. In: Proceedings of the 13th International Conference on Computer Communications and Networks (ICCCN). (2004) 373–378
12. Matoušek, J., Seidel, R., Welzl, E.: How to net a lot with little: Small ε-nets for disks and halfspaces. In: Proceedings of the Sixth Annual Symposium on Computational geometry (SCG). (1990) 16–22 Corrected version available at http://kam.mff.cuni.cz/∼matousek/oldpaps.html.
13. O'Rourke, J.: Art Gallery Theorems and Algorithms. Volume 3 of International Series of Monographs on Computer Science. Oxford University Press, New York, NY (1987)
14. Valtr, P.: Guarding galleries where no point sees a small area. Israel Journal of Mathematics **104** (1998) 1–16
15. Fusco, G., Gupta, H.: Selection and orientation of directional sensors for coverage maximization. In: Proceedings of 6th Annual IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON). (2009)
16. Alon, N., Spencer, J.H.: 13.4. In: The Probabilistic Method. 2nd edn. Wiley-Interscience (2000) 220–225

Appendix

A About the number of iterations of the doubling process

This appendix contains the proof that when M is equal to the size of the k-hitting set, then $\frac{k}{2M}$, $\frac{4}{k} \cdot M \log_2 \frac{n}{M}$ iterations of the doubling process are enough to retrieve the optimal k-hitting set. This proof follows the lines of the one in [6], but with the additional parameter k.

Theorem 4. *If $\varepsilon = \frac{k}{2M}$, $\frac{4}{k} \cdot M \log_2 \frac{n}{M}$ iterations of the internal* for *loop of Algorithm 1 are sufficient to find a k-hitting set.*

Proof. Initially $w(X) = |X| = n$. The set S selected at the end of the internal *for* loop satisfies $w(S) < \varepsilon \cdot w(X)$, because the algorithm found a weighted ε-net. Doubling the weights of the elements in S adds a total of $w(S)$ new weight to the system. So $w(X)$ grows at most by a $(1+\varepsilon)$ factor at each iteration. Then, after t iterations

$$w(X) \;\leq\; n\,(1+\varepsilon)^t$$

Let H^* be the optimal k-hitting set. Initially we have $w(H^*) = |H^*| = M$. Since H^* is a k-hitting set, there are least k elements of H^* in each set of $\mathcal{R}$. So for any possible set S chosen in step 11, there are at least k elements of H^* that are doubled. By the convexity of the function 2^x, the increase of $w(H^*)$ is minimal if the doublings are spread out over the elements of H^* as evenly as possible. So after t iterations, we have

$$w(H^*) \;\geq\; M \cdot 2^{kt/M}$$

Since the weights are positive and $H^* \subseteq X$, $w(H^*) \leq w(X)$. We need to find the largest t for which

$$M \cdot 2^{kt/M} \;\leq\; n \left(1 + \frac{k}{2\,M}\right)^t$$

can be true. Taking the log

$$\log_2 M + \frac{kt}{M} \;\leq\; \log_2 n + t\,\log_2\left(1 + \frac{k}{2\,M}\right)$$

and solving for t

$$t \;\leq\; \frac{\log_2 n - \log_2 M}{\frac{k}{M} - \log_2\left(1 + \frac{k}{2\,M}\right)} \;\leq\; \frac{4}{k} \cdot M \log_2 \frac{n}{M}$$

where we used the fact that $\log_2(1+x) \leq \frac{3}{2}x$ for $x > 0$. Since the expression on the RHS is $O(n)$ for any possible value of M, the theorem follows. □

B Computing Weighted (k, ε)-nets by Random Sampling

This appendix contains the proof of Theorem 2, which is an extension of the ε-net theorem of Haussler and Welz [7]. As explained in section 3.2, the two challenges in generalizing the random-sampling technique are that (i) sampling with replacement cannot be used, and (ii) weights must be part of the sampling process. Our contribution is a new method to obtain weighted (k, ε)-nets in which: (i) we sample a subset of points at once (without duplicates), and (ii) we include the weights directly in the sampling process.

We start by proving three lemmas, and then we will prove Theorem 2. Let m be as given by equation (1), and N be the subset of points randomly picked from X as described in Theorem 2. After picking N, pick another set T (for the purpose of the below analysis) in the same way as N. We now define two events

$$\begin{aligned} E_1 &= \{\exists A \in \mathcal{C} \text{ s.t. } |A \cap N| < k,\ w(A) \geq \varepsilon\, w(X)\} \\ E_2 &= \{\exists A \in \mathcal{C} \text{ s.t. } |A \cap N| < k,\ w(A) \geq \varepsilon\, w(X),\ |A \cap Z| \geq \varepsilon m\} \end{aligned}$$

where $Z = N \cup T$.

Intuitively, E_2 is the event that N does not k-hit some set $A \in \mathcal{C}$, but Z has a "large" intersection with the set A (also remember that T is disjoint from N). Note that $2\varepsilon m$ is a lower bound the average size of the intersection of A and Z (as computed below).

Lemma 1. $\Pr[E_2] \geq \frac{1}{2} \Pr[E_1]$.

Proof. $E_2 \subseteq E_1$, because if E_2 happens, then E_1 happens too. From the definition of conditional probability

$$\Pr[E_2 \mid E_1] = \frac{\Pr[E_2 \cap E_1]}{\Pr[E_1]} = \frac{\Pr[E_2]}{\Pr[E_1]}$$

so it suffices to show that $\Pr[E_2 \mid E_1] \geq 1/2$.

Let $Z = \{y_1, \ldots, y_{2m}\}$ (where the y_i's *are* pairwise different). Since E_1 happens, there is some set A s.t. $|A \cap N| < k$ and $w(A) \geq \varepsilon \cdot w(X)$. Therefore, $\Pr[E_2 \mid E_1]$ is at least the probability that, for this A, $|A \cap Z| \geq \varepsilon\, m$.

Let Y_i be the random variable (r.v.)

$$Y_i = \begin{cases} 1 & \text{if } y_i \in A \\ 0 & \text{o.w.} \end{cases}$$

Each subset N (resp. T) is picked with probability proportional to the sum of the weights of its elements, and each element can appear in $\binom{n-1}{m-1}$ (resp. $\binom{n-m-1}{m-1}$) subsets (because these are the ways of putting every other elements in the remaining positions). So the probability of picking one element depends only on its weight, and not on the other elements, which means that the elements are pairwise independent. So we have that

$$\Pr[Y_i = 1] = \frac{w(A)}{w(X)} \geq \frac{\varepsilon \cdot w(X)}{w(X)} = \varepsilon$$

where $w(\cdot)$ is a function that returns the weight of given set, which is defined as the sum of the weights of its elements.

Let $Y = \sum_{i=1}^{2m} Y_i$. Clearly $Y = |A \cap Z|$. We have

$$\mathrm{E}[Y] = \mathrm{E}[\sum_{i=0}^{2m} Y_i] = \sum_{i=0}^{2m} \mathrm{E}[Y_i] = \sum_{i=0}^{2m} \Pr[Y_i = 1] = \sum_{i=0}^{2m} \frac{w(A)}{w(X)} = 2m \frac{w(A)}{w(X)} \geq 2\,\varepsilon\, m$$

To bound the deviation from the expectation, we use Chebyshev's inequality:

$$\Pr[|Y - \mathrm{E}[Y]| > t] \leq \frac{\mathrm{Var}[Y]}{t^2}$$

Since Y_i's are independent, the covariance is zero. So we get

$$\begin{aligned}\mathrm{Var}[Y] &= \sum_{i=0}^{2m} \mathrm{Var}[Y_i] \;=\; 2m\,\mathrm{Var}[Y_i] \\ &= 2m(\mathrm{E}[Y_i^2] - (E[Y_i])^2) \\ &= 2m\left(\frac{w(A)}{w(X)} - \left(\frac{w(A)}{w(X)}\right)^2\right) \\ &\le 2m\,\frac{w(A)}{w(X)}\end{aligned}$$

where we used the fact that for the 0-1 r.v. Y_i, $Y_i = Y_i^2$.

Applying Chebyshev's inequality

$$\begin{aligned}\Pr[Y < \varepsilon\, m] &\le \Pr[Y < m\frac{w(A)}{w(X)}] \\ &= \Pr[Y < 2m\frac{w(A)}{w(X)} - m\frac{w(A)}{w(X)}] \\ &\le \Pr[|Y - \mathrm{E}[Y]| > m\frac{w(A)}{w(X)}] \\ &\le \left(\frac{1}{m}\frac{w(X)}{w(A)}\right)^2 \mathrm{Var}[Y] \\ &\le \frac{4}{m}\frac{w(X)}{w(A)} \;\le\; \frac{4}{\varepsilon\, m} \\ &\le \frac{1}{2}\end{aligned}$$

where in the last inequality we used the fact that

$$m \;\ge\; 4(d+k-1)/\varepsilon \log_2(4(d+k-1)/\varepsilon) \;\ge\; 8/\varepsilon$$

Finally we have that

$$\Pr[E_2 \mid E_1] \;\ge\; \Pr[Y \ge \varepsilon\, m] \;\ge\; \frac{1}{2}$$

□

Lemma 2. $\Pr[E_2] \le g(d, 2m)\,(2m)^{2k-2}\,2^{-\varepsilon m}$,
where $g(d, 2m) = \binom{n}{0} + \binom{n}{1} + \ldots + \binom{n}{d} \le n^d$.

Proof. The experiment of picking N and T can be viewed in an alternative way. Pick a subset $Z \subseteq X$ of size $2m$ at random (each subset is picked with probability proportional to the sum of the weights of its elements). Then, pick N as a subset of Z of size m at random (again with probability proportional to the sum of the weights of its elements). Finally, let $T = Z \setminus N$. Note that this view is equivalent because the probability of picking any subset N is the same as before, (similarly for T). This can be verified as follow. We are going to compute the probability of picking a certain subset $\overline{N}$ in both cases. In order to do this, we need to compute the sum of all possible sets of size m. Among all possible sets of size m, each element appears in exactly $\binom{n-1}{m-1}$ of them (because these are the ways of putting every other element in the remaining positions). Now, it is not necessary to know exactly in which set each element gives its contribution, but it is enough to know that it appears a total of $\binom{n-1}{m-1}$ times. So, the sum of weight of all possible sets of size m is $\binom{n-1}{m-1}w(X)$. Then

$$\begin{aligned}\Pr[\text{picking } \overline{N} \text{ directly from } X] &= \frac{\text{weight of } \overline{N}}{\sum \text{weight of all possible subsets of size } m} \\ &= \frac{w(\overline{N})}{\binom{n-1}{m-1}w(X)}\end{aligned}$$

Now we are going to compute the probability of picking a subset Z containing $\overline{N}$. This requires to determine the sum of the weights of all subsets of size $2m$ that contain $\overline{N}$. $\overline{N}$ appears in $\binom{n-m}{m}$ of them (as these are the number of ways of putting any other element in the remaining m positions), and it gives a contribution of $w(\overline{N})$ in each of them. Any other element can appear in any of the remaining m positions, for a total of $\binom{n-m-1}{m-1}$ times (because fixed any one them, these are the ways of putting any of the remaining $n-m-1$ in the other $m-1$ positions). So we get that

$$\begin{aligned}\Pr[\text{picking } Z \text{ containing } \overline{N}] &= \frac{\sum \text{weight of all subsets of size } 2m \text{ containing } \overline{N}}{\sum \text{weight of all possible subsets of size } 2m} \\ &= \frac{\binom{n-m}{m} w(\overline{N}) + \binom{n-m-1}{m-1} w(X \setminus \overline{N})}{\binom{n-1}{2m-1} w(X)}\end{aligned}$$

We also need to compute the probability of picking $\overline{N}$ from Z

$$\begin{aligned}\Pr[\text{picking } \overline{N} \text{ from } Z \mid \overline{N} \subset Z] &= \frac{\text{weight of } \overline{N}}{\sum \text{weight of all possible subsets of } Z \text{ of size } m} \\ &= \frac{w(\overline{N})}{\binom{2m-1}{m-1} w(Z)}\end{aligned}$$

This requires to know $w(Z)$, which can be computed as

$$\begin{aligned}w(Z) &= \frac{\sum \text{weight of all subsets of size } 2m \text{ containing } \overline{N}}{\text{number of ways of choosing } \overline{N}} \\ &= \frac{\binom{n-m}{m} w(\overline{N}) + \binom{n-m-1}{m-1} w(X \setminus \overline{N})}{\binom{n-m}{m}}\end{aligned}$$

So

$$\Pr[\text{picking } \overline{N} \text{ from } Z \mid \overline{N} \subset Z] = \frac{\binom{n-m}{m} w(\overline{N})}{\binom{2m-1}{m-1} \left(\binom{n-m}{m} w(\overline{N}) + \binom{n-m-1}{m-1} w(X \setminus \overline{N}) \right)}$$

Finally, it is easy to verify that

$$\begin{aligned}\Pr[\text{picking } \overline{N} \text{ directly from } X] &= \Pr[\text{picking } Z \text{ containing } \overline{N}] \\ &\quad \cdot \Pr[\text{picking } \overline{N} \text{ from } Z \mid \overline{N} \subset Z]\end{aligned}$$

Let $A \in \mathcal{C}$, with $w(A) \geq \varepsilon \cdot w(X)$, and define $E_A = \{|A \cap N| < k,\ |A \cap Z| \geq \varepsilon\, m\}$.

Since N and T are disjoint, $|A \cap Z| = |A \cap N| + |A \cap T|$, and then $|A \cap N| < k$ is equivalent to $|A \cap Z| < k + |A \cap T|$. If $|A \cap Z| < \varepsilon m$, then E_A does not happen, and it doesn't happen if $|A \cap Z| > m + k - 1$ either.

Suppose that $|A \cap Z| = \varepsilon m + j$, where $0 \leq j \leq m + k - 1 - \varepsilon m$, then we can pick N as follow. We select $m - k + 1$ elements among the points outside the intersection with A

$$\binom{2m - \varepsilon m - j}{m - k + 1}$$

and the remaining $k - 1$ elements anywhere else

$$\binom{m + k - 1}{k - 1}$$

Their product can be bounded in the following way

$$\begin{aligned}
\binom{2m-\varepsilon m-j}{m-k+1}\binom{m+k-1}{k-1} &\leq \binom{2m-\varepsilon m}{m-k+1}\binom{m+k-1}{k-1} \\
&= \frac{(2m-\varepsilon m)!}{(m-k+1)!(m-\varepsilon m+k-1)!}\,\frac{(m+k-1)!}{(k-1)!m!} \\
&= \frac{(2m-\varepsilon m)!}{m!(m-\varepsilon m)!}\,\frac{(m+k-1)!}{(m-k+1)!}\,\frac{(m-\varepsilon m)!}{(m-\varepsilon m+k-1)!}\,\frac{1}{(k-1)!} \\
&\leq (2m)^{2k-2}\binom{2m-\varepsilon m}{m}
\end{aligned}$$

where in the last inequality we used the fact that

$$\frac{(m+k-1)!}{(m-k+1)!}\,\frac{(m-\varepsilon m)!}{(m-\varepsilon m+k-1)!}\,\frac{1}{(k-1)!} \leq (2m)^{2k-2}$$

which can be proved by induction. The base case, $k=1$, is trivial. Assuming that the formula is valid for $k-1$, we get

$$\begin{aligned}
\frac{(m+k)!}{(m-k)!}\,\frac{(m-\varepsilon m)!}{(m-\varepsilon m+k)!}\,\frac{1}{k!} &= \frac{(m+k)\,(m+k-1)!}{(m-k+1)!\,/\,(m-k+1)}\,\cdot \\
&\quad \cdot\frac{(m-\varepsilon m)!}{(m-\varepsilon m+k)\,(m-\varepsilon m+k-1)!}\,\frac{1}{k\,(k-1)!} \\
&\leq \frac{(m+k)\,(m-k+1)}{k\,(m-\varepsilon m+k)}\,(2m)^{2k-2} \\
&\leq \frac{(m+1+k)(m+1-k)}{k\,(m-m+k)}\,(2m)^{2k-2} \\
&\leq (m+1)^2\,(2m)^{2k-2} \\
&\leq (2m)^{2k}
\end{aligned}$$

where in the last inequality we used the fact that $m \geq 1$.

Using this fact, we can bound $\Pr[E_A]$. Recall that E_A happens only if $\varepsilon m \leq |A \cap Z| \leq m+k-1$. So

$$\begin{aligned}
\Pr[E_A] &= \Pr[|A\cap N| < k,\ |A\cap Z| \geq \varepsilon m] \\
&\leq \Pr[N \text{ chosen s.t. } |A\cap N| < k \ \mid\ \varepsilon m \leq |A\cap Z| \leq m+k-1] \\
&\leq (2m)^{2k-2}\,\frac{\binom{2m-\varepsilon m}{m}}{\binom{2m}{m}} \\
&= (2m)^{2k-2}\,\frac{(2m-\varepsilon m)\cdots(m+1)\cdot m\cdots(m-\varepsilon m+1)}{2m\cdots(2m-\varepsilon m+1)\cdot(2m-\varepsilon m)\cdots(m+1)} \\
&= (2m)^{2k-2}\,\frac{m\cdots(m-\varepsilon m+1)}{2m\cdots(2m-\varepsilon m+1)} \\
&\leq (2m)^{2k-2}\,2^{-\varepsilon m}
\end{aligned}$$

For two sets $A, A' \in \mathcal{C}$ s.t. $w(A),\, w(A') \geq \varepsilon\, w(X)$ and $A \cap Z = A' \cap Z$, the events E_A and $E_{A'}$ are the same. This is because the occurrence of E_A depends only on the intersection $A \cap Z$. The number of sets

$A \in \mathcal{C}$ s.t. $A \cap Z$ is unique is at most $|\{A \cap Z \mid A \in \mathcal{C}\}| = |\mathcal{C}_{|Z}| \leq g(d, 2m)$ by Corollary 1 below. Then

$$\begin{aligned} \Pr[E_2] &\leq \bigcup_{A \mid A \cap Z \text{ is unique}} \Pr[E_A] \\ &\leq |\mathcal{C}_{|Z|}|\,(2m)^{2k-2}\,2^{-\varepsilon m} \\ &\leq g(d, 2m)\,(2m)^{2k-2}\,2^{-\varepsilon m} \end{aligned}$$

□

Lemma 3. *For any set system $(X, \mathcal{C})$, with $|X| = n$ and VC-dimension d, $|\mathcal{C}| \leq g(d, n)$, where $g(d, n) = \binom{n}{0} + \binom{n}{1} + \ldots + \binom{n}{d} \leq n^d$.*

Proof. See [16]. □

Given a set system $(X, \mathcal{C})$, for any subset of the points $N \subseteq X$, let $\mathcal{C}_{|N}$ denote the projection of $\mathcal{C}$ onto N, i.e. the set $\{A \cap N \mid A \in \mathcal{C}\}$.

Corollary 1. *For any set system $(X, \mathcal{C})$, if $N \subseteq X$, then $(N, \mathcal{C}_{|N})$ has VC-dimension $\leq d$, which implies $|\mathcal{C}_{|N}| \leq |N|^d$.*

Finally, we prove the main theorem.

Proof of theorem 2. Combining Lemmas 1, 2 and 3

$$\Pr[E_1] \leq 2\Pr[E_2] \leq 2g(d, 2m)m^{2k-2}2^{-\varepsilon m} \leq 2(2m)^{d+2k-2}2^{-\varepsilon m}$$

So we need to show that

$$2\,(2m)^{d+2k-2}\,2^{-\varepsilon m} \leq \delta$$

which can be written as

$$\frac{2}{\delta}(2m)^{d+2k-2} \leq 2^{\varepsilon m}$$

or equivalently

$$\varepsilon m \geq \log\frac{2}{\delta} + (d + 2k - 2)\log(2m)$$

Now we consider each part of the sum separately. From (1), it follows that

$$\frac{1}{2}\varepsilon m \geq \log\frac{2}{\delta}$$

so it suffice to show that

$$\frac{1}{2}\varepsilon m \geq (d + 2k - 2)\log(2m)$$

If this inequality is valid for some value of m, then it is for any valid for any bigger value of m. So we just need to verify it for $m = 4(d + 2k - 2)/\varepsilon \log(4(d + 2k - 2)/\varepsilon)$. Plugging in m we get

$$\begin{aligned} 2(d + 2k - 2)\log\frac{4(d + 2k - 2)}{\varepsilon} &\geq \\ &\geq (d + 2k - 2)\log\left(\frac{8(d + 2k - 2)}{\varepsilon}\log\frac{4(d + 2k - 2)}{\varepsilon}\right) \end{aligned}$$

and is equivalent to

$$\frac{2(d + 2k - 2)}{\varepsilon} \geq \log\frac{4(d + 2k - 2)}{\varepsilon}$$

which is definitely true. □

C Computing *Small* Weighted (k, ε)-nets for disks

In this appendix we present a simple extension of [12] to build *small* weighted (k, ε)-nets for disks. The original construction in [12] easily extends to (k, ε)-nets by replacing $\delta = \varepsilon/6$ with $\delta = \varepsilon/(6k)$. The proof presented here is simplified respect to the original one, because we consider only disks, instead of pseudo disks.

The underlining idea is to pick points that are spaced apart, which hit all large enough disks. The strategy that we are going to use is to draw "colored" disks that contain a fixed number of points, and select points only on the border of the colored disks. The position and the size of colored disks depend on the input points, but not on the input disks. All colored disks, but one, will have exactly $\lfloor \delta n \rfloor$ input points, where $\delta = \varepsilon/(6k)$. Each input point gets the color of the colored disks that covers it, or remains uncolored if uncovered. After placing the colored disks, we compute a Dealaunay triangulation (DT) of the colored points. DT will have uni-colored, bi-colored, and tri-colored triangles. Triangles will have uni-colored and bi-colored edges. Let's define some terminology (see Figure 3):

Definition 7 (Corridor; Hall; Sides; Ends; Corners).

- *Let a* corridor *be a maximal connected chain of bi-colored triangles in DT sharing bi-colored edges. In our construction, corridors are between two colored disks.*
- *Let a* hall *be a maximal group of adjacent tri-colored triangles (this is a generalization of the* degenerate-corridors *of [12]). In our construction, halls are between 3 or more disks, attached to the end of the corridors.*
- *Corridors are bounded by two chains of uni-colored edges, which we call* sides*, and two bi-colored edges, which we call* ends.
- *We call the endpoints of the sides the* corners *of the subcorridor. Note that one of the sides can degenerate in a single point, in which case there are 3 corners, instead of 4.*

We start by describing the algorithm for the unweighted case, and we will show how to add weights afterwords. We are given $0 < \varepsilon \leq 1$, a family $\mathcal{D}$ of disks, and a set X of $n > 2$ points in the plane. For simplicity assume that the points are in general position, (i.e. no three points are collinear, and no four points are cocircular). Let define $\delta = \varepsilon/(6k)$ (the reason for this will be clear soon). Let $X_1, \ldots, X_j$ be disjoints subsets of X constructed in the following manner. From the boundary of X, "bite off" subsets of X of size $\lfloor \delta n \rfloor$ with the following properties:

- The union of all the X_i subsets, contains the boundary points of X: $conv(X) \subseteq \bigcup_{1 \leq i \leq j} X_i$ (where *conv* is the *convex hull*)
- each X_i is representable as $X \cap H_i$ for some halfplane H_i (or equivalently, $X_i = X \cap D_i$ for some (large enough) disk D_i, and to simplify the proof we can think that D_i is bigger than any input disk)
- $|X_i| = \lfloor \delta n \rfloor$ for $1 \leq i < j$, and $|X_j| \leq \lfloor \delta n \rfloor$

Now, consider the internal points of X, that are not part of any disk X_i. We are going to draw the largest number of disks of size $\lfloor \delta n \rfloor$ to cover the internal points. Specifically, let $X_{i+1}, \ldots, X_r$ be a maximal collection of disjoint subsets of $X \setminus \bigcup_{1 \leq i \leq j} X_i$ satisfying:

- $X_i = X \cap D_i$ for some disk D_i
- $|X_i| = \lfloor \delta n \rfloor$ for $j < i \leq r$.

At this point, we have a total of r disks X_i. For each i, $1 \leq i \leq r$, color the points of X_i with color i, and call D_i the *disk defining color i*, or *colored disk i*. Let $\overline{X}$ be the set of colored points, and call the points in $X \setminus \overline{X}$ *colorless*. Let DT be the Dealaunay triangulation of the set of colored points $\overline{X}$. Break each corridor C into a minim number of *subcorridors*, i.e. subchains of the chain of triangles that form C, so that each subcorridor contains at most $\lfloor \delta n \rfloor$ colorless points. Let N be the set of the corners of all subcorridors. Clearly $N \subset X$. We are going to show that N is a the (k, ε)-net for $\mathcal{D}$, i.e. any disk of $\mathcal{D}$ that contains εn points of X, also contains k points of N. This construction is summarized in Algorithm 2.

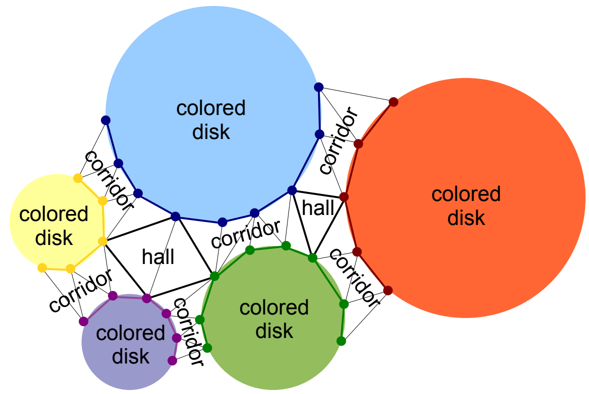

Fig. 3. Colored disks, corridors, and halls.
Note that points are not necessarily on the boundaries of the disks, but they are drawn this way to make the picture clearer

First of all note that colorless points can only be in corridors and halls (because unicolored triangles are contained in the corresponding color-defining disks). Also, we can observe that any disk $D \in \mathcal{D}$ containing no colored points, contains less than $\lfloor \delta n \rfloor$ points of X. In fact, from the maximality of the construction, there cannot be colorless disks with $\lfloor \delta n \rfloor$ points. Than we claim that:

Claim. There are at most $3r - 6$ corridors in DT, and $r \leq \lceil 1/\delta \rceil + 1$

Proof. See [12]. Note that, since all X_i are disjoint, and all but maybe one contain $\lfloor \delta n \rfloor$ points, so $r \leq \lceil 1/\delta \rceil + 1$. □

We now prove the (k, ε)-net theorem for disks.

Theorem 5 ((k, ε)-net theorem for disks). *Algorithm 2 creates a (k, ε)-net of size $O(k/\varepsilon)$ for $(X, \mathcal{D})$, where X is a set of points ($|X| > 2$) in non-$\mathcal{D}$-degenerate position (i.e. no three points are collinear, and no four points are cocircular), and $\mathcal{D}$ is a family of disks.*

Proof. By claim C we have that the size of N is $O(k/\varepsilon)$. So we only need to show that N contains at least k points in each input disk of size at least εn.

First of all, the case case $k = 1$ is already proved in [12], so we focus on the case of $k \geq 2$.

For a generic input disk of size at least εn, we are going to compute the minimum number of corners contributed by each intersecting region (colored disk, subcorridor or hall), while assuming that it has the largest possible intersection. Also, we will pay attention not to count the same corner multiple times. If a colored disk is completely contained in an input disk, it will contribute for the biggest number of corners. So we should only consider colored disks that intersect the boundary of the input disk. We claim that the minimum contribution is given when the boundary of an input disk intersect an alternation of colored disks and corridors. In this case we should count 1 corner, for each colored disk/corridor. The only case in which (a part of) the boundary of the input disk does not intersect any colored disk is when it is contained inside a corridor. But this can only happen if there is a colored disk inside the input disk, but we already argued that this will give a higher contribution. The following is an upper bound on the number of intersecting points. There can be δn points for the colored disk, plus another δn points for the subcorridor on the side of the corner that we are counting, plus another δn points for the hall adjacent to them. This means that for $3\delta n$ points there is at least 1 corner. We are considering input disks of size at least $\varepsilon n = 6k\delta n$, and this implies that there are at least k points in each disk. □

Algorithm 2: Small (k, ε)-nets for disks

1 Let $\delta = \varepsilon/(6k)$;

2 Let $X_1, \ldots, X_j$ be disjoints subsets of X with the following properties:

- $conv(X) \subseteq \bigcup_{1 \leq i \leq j} X_i$ (where *conv* is the *convex hull*)
- each X_i is representable as $X \cap H_i$ for some halfplane H_i (or equivalently, $X_i = X \cap D_i$ for some (large enough) disk D_i, and to simplify the proof we can think that D_i is bigger than any input disk)
- $|X_i| = \lfloor \delta n \rfloor$ for $1 \leq i < j$, and $|X_j| \leq \lfloor \delta n \rfloor$

This construction of subsets of X can be done by repeatedly "biting off" subsets of X with halfplanes;

3 Let $X_{i+1}, \ldots, X_r$ be a maximal collection of disjoint subsets of $X \setminus \bigcup_{1 \leq i \leq j} X_i$ satisfying:

- $X_i = X \cap D_i$ for some disk D_i
- $|X_i| = \lfloor \delta n \rfloor$ for $j < i \leq r$;

4 For each i, $1 \leq i \leq r$, color the points of X_i with color i, and call D_i the *disk defining color* i, or *colored disk* i. Let $\overline{X}$ be the set of colored points, and call the points in $X \setminus \overline{X}$ *colorless*;

5 Let DT be the Dealaunay triangulation of the set of colored points $\overline{X}$;

6 Break each corridor C into a minim number of *subcorridors*, i.e. subchains of the chain of triangles that form C, so that each subcorridor contains at most $\lfloor \delta n \rfloor$ colorless points;

7 Let N be the set of the corners of all subcorridors. Clearly $N \subset X$;

8 Return N has the (k, ε)-net for $\mathcal{D}$;

Finally, we consider the weighted case. The construction is similar to the unweighted one, with the only difference that the colored disks contain $\lfloor \delta w(X) \rfloor$ points, where $w(X)$ is the sum of the weights of all points. It is easy to see that the proof follows through.